

# Towards Moral Autonomous Systems\*

Vicky Charisi, Louise Dennis, Michael Fisher,  
Robert Lieck, Andreas Matthias, Marija Slavkovik,  
Janina Sombetzki, Alan F. T. Winfield, Roman Yampolskiy

March 20, 2017

## Abstract

Both the ethics of autonomous systems and the problems of their technical implementation have by now been studied in some detail. Less attention has been given to the areas in which these two separate concerns meet. This paper, written by both philosophers and engineers of autonomous systems, addresses a number of issues in machine ethics that are located at precisely the intersection between ethics and engineering. We first discuss different approaches towards the conceptual design of autonomous systems and their implications on the ethics implementation in such systems. Then we examine problematic areas regarding the specification and verification of ethical behavior in autonomous systems, particularly with a view towards the requirements of future legislation. We discuss transparency and accountability issues that will be crucial for any future wide deployment of autonomous systems in society. Finally we consider the, often overlooked, possibility of intentional misuse of AI systems and the possible dangers arising out of deliberately unethical design, implementation, and use of autonomous robots.

## 1 Introduction

The so called “trolley problem” is a thought experiment introduced in [22], whose ethical conundrum continues to fascinate today. Roughly, it can be summarized as follows: there is a runaway trolley on a railroad track and ahead on the track there are five people tied up, unable to escape being killed by the trolley. The track splits in two by a lever you control. The lever can divert the trolley on to a second track where there is one person tied up and unable to move. Is it more ethical to divert the train, or let it run its course?

The emergence of driver-less cars in regular traffic has brought the trolley problem to public attention. Articles such as “Should Your Car Kill You To Save Others?”<sup>1</sup> are flooding popular science media. It is easy, given the same

---

\*This article is the result of a series of discussions taken within the scope of the Dagstuhl Seminar 6222 [19]. Dennis, Fisher and Winfield wish to thank EPSRC for their support, particularly via the “Verifiable Autonomy” research project (EP/L024845 and EP/L024861) Corresponding Author: Andreas Matthias, Philosophy Department, Lingnan University, Tuen Mun, N.T. Hong Kong, matthias@ln.edu.hk; Corresponding Author: Marija Slavkovik, University of Bergen, P.O.Box 7802, 5020 Bergen, Norway, marija.slavkovik@uib.no

<sup>1</sup><http://www.popularmechanics.com/cars/a21492/the-self-driving-dilemma/>

problem context of traffic, to get sidetracked into confusing solving the trolley problem with controlling the impact driver-less cars will have on traffic and our society as a whole, but this is not the case. Enabling machines to exhibit ethical behavior is a very complex and very real time sensitive issue. The driver-less cars are only the forefront of a whole generation of intelligent systems, machines and software, that can operate autonomously and will operate as part of our society. Which ethical theory to employ for deciding whose death to avoid in a difficult traffic situation is a difficult problem, but not necessarily the most important one we need to solve. The goal of this position paper is to discuss what does *engineering machine ethics* entail, namely once we as a society have discerned what is the right thing to do, how can we make sure the machine does it?

The problem of identifying, discerning, and recommending concepts of right and wrong is the domain of moral philosophy. Moral philosophy, together with the law, act as a system of recommendations regarding which possible actions are to be considered right or wrong (ignoring, for the moment, ethics systems that do not specifically address the morality of individual *actions*, e.g. character-based ethics, and which are less useful for the problem at hand).

Driver-less vehicles are the most visible of a whole range of technologies including also assisted living technologies and various embedded decision aid software and solutions. We have been using intelligent systems with a varying degree of autonomy for the past fifty years: industrial robots, intelligent programming of household appliances, automated trains, etc. However, what all of these machines have in common is that they either operate in a segregated space, a so called *working envelope*, or they have no capabilities to cause damage to their environment. Driver-less vehicles are obviously going to “break” both these restrictions. How can we build intelligent autonomous systems that uphold the ethical values of the society in which they are embedded? This is the main concern of *machine ethics*, a new interdisciplinary area of research within Artificial Intelligence (AI) [36, 1, 55, 4].

It is clear that choosing the best moral theory to implement in a particular intelligent autonomous system is not a simple question. It is primarily a question for moral philosophy and opens new challenges in this field. For a long time, people and societies have been the only intelligent decision-makers and moral philosophy has been developed with the, often implicit, assumption that the moral agent is a human. It is not at all clear to what extent existing moral theories extend to non-human decision-makers. Even if it is shown to be easy to replace a human agent with an artificial agent in a moral theory, and some societal decision is made concerning which ethical behavior in machines is desirable or sufficient, we are still faced with a set of problems regarding the *implementation* of ethical reasoning. These are the problems we analyze here: implementation, verification, trust, confidence and transparency and the prevention of intentionally unethical systems.

Moral theories are inherently ambiguous in recommendations of moral behavior, thus reflecting the context dependency of what constitutes a moral choice, and we already know that artificial systems, unlike people, are not good at handling ambiguity. Enabling machines to deal with context-ambiguity in decision-making is a core Artificial Intelligence problem [43]. In Section 2 we give an overview of the most intuitive approaches to implementing ethical behavior in autonomous systems and discuss the advantages and shortcomings of

these approaches.

Human societies have a multitude of means for ensuring its members behave within the socially accepted boundaries of morality. We can say that a person behaves ethically because they have an individual and personal motivation to do so, without going into how this motivation is formed. The motivation for an artificial agent to behave ethically originates not personally from the agent, but from other actors. These actors can broadly be organized into three groups: the manufacturers of the artificial agent, its users, and the various societal regulators whose job it is to make sure that order in society is maintained. It is all of these actors that need to be reassured that their own particular motivations for the AI system to behave ethically are met. Hence, a big concern when engineering machine ethics is not only that ethical behavior is accomplished, but also that the ethical behavior can be verified. This is the issue we discuss in Section 3.

In Section 4 we focus on issues of transparency and accountability of machine ethics implementations. Since there are several actors outside of the ethical agent who are supplying the motivation for ethical behavior, the implementation of this behavior must be transparent to those actors to the extent and in a manner sufficient for their needs. Transparency is a key element in enabling society to have the right amount of trust and confidence in the operations of an AI system.

Lastly in Section 5 we discuss the possible dangers for society that lie in engineering machine ethics. Like all technology, AI systems can also be abused to further criminal activities. AI systems can be deliberately built to behave unethically and illegally, but they also can be induced, deliberately or by accident, into exhibiting socially undesirable behavior.

The main contribution of this position paper is an integral overview of the immediate challenges and open questions faced when pursuing the problem of engineering of machine ethics. This paper is a result of a week long discussion among experts of different fields within the scope of the Dagstuhl Seminar 16222<sup>2</sup>, and incorporates various ideas that arose as a result of discussions among interdisciplinary experts. Position papers that focus on machine ethics as a whole have been produced and they offer interesting insights in the problem as a whole, see for example [36, 3, 8, 4], but to the best of our knowledge, this is the only document devoted specifically to the problem of engineering machine ethics.

## 2 Different approaches, their advantages and challenges

An intelligent system is one that is capable of communicating with, and reasoning about, its environment and other systems. An autonomous system is one that is capable of, to a certain extent, unsupervised operation and decision-making. Artificial Intelligence (AI) has grown into a large field that incorporates many approaches, which can be, very tentatively, classified into *soft computing approaches*, which include statistical methods, machine learning and probabilistic reasoning, and *traditional symbolic AI methods*, which includes logic-based reasoning [43]. The question of how to implement machine ethics in an intelli-

---

<sup>2</sup>The report on this seminar is available [19].

gent autonomous system necessarily hinges on the AI methods that system uses. Different AI approaches would be subject to different machine ethics implementations and we need to consider their malleability to machine ethic approaches, as well as their risks and advantages with this respect.

Very intuitively, ethical behavior in machines can be accomplished in at least two different ways: by implementing moral agency or by constraining unethical behavior. Wallach, Allen, and Smith refer to these two approaches as the *top-down* and *bottom-up* approach respectively [55], and we will use their terminology throughout the paper. A moral agent is an agent capable of discerning right from wrong. Moral behavior is obtained, without imbuing the AI system with moral agency, when the unethical actions, decisions, or options in general, are recognized as such by a person and made unavailable to the AI system. Figure 1 illustrates the Top-down and bottom-up approach for building an AI system. We discuss both of these approaches and their advantages and challenges.




Figure 1: Top-down and bottom-up approach for building an AI system. **(left):** Behavior emerges from the interplay of the agent, the task it is to perform, and the environment. **(right):** From a more technical point of view the agent is defined as an algorithm (and its embodiment in case of physical agents), the task is formulated as an objective, and the environment generates data that are processed by the agent.

## 2.1 What are the top-down and bottom-up approaches

Within engineering, a top-down approach towards solving a task consists in breaking down the task iteratively into smaller sub-tasks until one obtains tasks that can be directly implemented. Within machine ethics, a top-down approach consists in taking a specified ethical theory and instantiating it to identifying particular states and actions as ethical or unethical with respect to that theory [55].

An example of an ethical reasoning approach that implements the top-down approach can be found in [14]. This paper considers a hybrid autonomous system three-layer architecture: a continuous system controlled by a rational software agent, which makes discrete decisions, via a continuous control layer that allows for a dynamic environment interaction and feedback. The rational software agent is provided with an ethical policy, a total order over abstract ethical principles such as “do no harm”, “do not damage property” etc. The agent relies on external entities to identify if, and which, of her possible actions impinges on some of these abstract ethical principles. Having her actions annotated, the

agent chooses between possible actions by selecting the one that is minimally unethical with respect to the given ethical policy.

An example of a bottom-up approach to machine ethics can be found in [2]. Here a system is developed that uses inductive logic programming over a corpus of particular cases of ethical dilemma to discover ethical preference principles. Each case relates two actions, one more ethical than the other. The preference between the actions depends on ethically relevant features that actions involve such as harm, benefit, respect for autonomy, etc. Each feature is represented as an integer that specifies the degree of its presence (positive value) or absence (negative value) in a given action. The system is able to extract an ethical rule from the cases it is presented with and thus to a certain extent is able to learn to discern right from wrong.

Top-down and bottom-up are two major approaches for developing an AI system. The two approaches tackle the problem of building an AI system from different sides and are not mutually exclusive. Each of the approaches comes with its own advantages and challenges with respect to building ethical AI.

## 2.2 Advantages and challenges of the top-down approach

The top-down approach is best suited for ethical reasoning in limited domains when the context of the decision-making problems can be predicted. Under the top-down approach, the ethical principles can be decided on before deploying the system and do not change throughout its usage. This allows for a thorough and well-informed process of decision-making and the verification of the system prior its practical application. Different ethical principles and theories may explicitly be encoded into the system giving clear options to decide upon. Any parameters that are left open for definition by the customer or to be learned from interaction with the environment have a clear function and it is possible to verify that they do not violate more general ethical principles.

The top-down approach also comes with its set of challenges. General ethical principles are typically formulated on a very abstract level. Much philosophical discourse on ethics is concerned with problems occurring when applying such general principles to concrete situations. It is thus not clear if, and how, general ethical principles can be leveraged to solve concrete problems of decision making in practice. For a real-world system the connection to the non-discrete sensory-motor level must be made. There are many ways to transform continuous sensor values into discrete propositions that can be used in reasoning. General principles or even single terms and concepts are only (if at all) implementable in a reduced way, i.e. restricted to one preferably clear interpretation. Due to their context sensitive definition it is not possible to consider every possible interpretation of an abstract principle or term in implementing them in an artificial system [32].

Furthermore, top-down approaches risk conflicts between the implemented ethical theories and principles. If only one theory is implemented, e.g. Kant's Categorical Imperative [18] or Isaac Asimov's first Law of Robotics<sup>3</sup>, then this theory would determine the specific maxims that are to be defined situationally by the artificial system. [58] describes experimental trials of a minimally ethical

---

<sup>3</sup>The laws can be found quoted in Wikipedia, at [http://en.wikipedia.org/wiki/Three\\_Laws\\_of\\_Robotics](http://en.wikipedia.org/wiki/Three_Laws_of_Robotics)

robot which implements Asimov’s three laws of robotics. The chosen theory must be such as to allow implementable rules to be derived from it. Such a monistic approach assumes that there exist no moral dilemmas, i.e. that the implemented theory is able to give a conflict-free rule to make a decision in every context ([58] experimentally shows how a single ethical rule performs when faced with a balanced ethical dilemma). What is missing in every top-down approach is what some might call “common sense”, creativity, or sensibility for phenomena in a world that we cannot completely predict.

Deciding on a specific set of ethical principles involves settling long-standing philosophical disputes in an ad-hoc way. It is possible that governmental bodies might take the lead in outlining high-level ethical guidance to manufacturers. For example, Germany’s ministry of transport recently announced the intention to set out a basic ethical policy to be followed by car manufacturers stipulating that property damage takes always precedence over personal injury, that there must be no classification of people, for example, on size, age and the like, and that – ultimately – it is the manufacturer who is liable<sup>4</sup>.

### 2.3 Advantages and challenges of the bottom-up approach

Under a bottom-up approach we know the kind of data the AI system receives from interacting with the environment and, while the overall objective might not be well known or specified, we still have an idea how to process these data in a useful manner. For instance, the vision pipeline of a household robot is designed to extract obstacles (walls, tables, etc.), objects of interest (books on a shelf etc.), and its own position from the sensory data because this information is useful for a wide range of tasks it will be required to perform. Most real-world AI systems will be partly designed in a bottom-up fashion at least on the lower sensory-motor level. With respect to the task of building ethical AI two major questions are: How can ethical principles be encoded in a bottom-up fashion? and, even if the ethical principles themselves do not follow the bottom-up approach, How do components designed in a bottom-up fashion affect the overall ethical properties of the system?

Bottom-up approaches do not predetermine moral principles, ethical theories or sets of rules, but instead formulate basal parameters and intend to implement competences whereby an artificial system acts autonomously. This can be done, for example, via trial and error or other modes of learning such as imitation, induction and deduction, exploration, learning through reward, association and conditioning [11]. Bottom-up approaches can be separated into models of evolution [23] and models of human socialization [21, 9]. The former simulate evolutionary moral learning, by assessing slightly different programs in an artificial system to evaluate an ethical case. Those programs that can solve the ethical task sufficiently go through to a “next round” where they are (re)combined to solve further ethical tasks. Evolutionary approaches can be used in earlier stages of moral development before considering models of human socialization.

Models of human socialization consider the role of empathy and emotion for moral learning. They assume that a robot learns morality via empathy

<sup>4</sup><http://www.wiwo.de/politik/europa/selbstfahrende-autos-dobrindt-gruendet-ethikkommission-fuer-automatisiertes-fahren/14513384.html>

[47]. What is controversial in the philosophical discourse is that there exist two types of empathy [51]: perceptual empathy, when an emotion triggers an equivalent or congruent reaction in the observer [34], and imaginative empathy that requires a change in perspective in the form of empathising with the other, putting oneself in the observed other’s shoes. Perceptual empathy is explicable with the help of specific *theories of mind* or neuronal resonance and mirror neurons and has been implemented in a rudimentary fashion in artificial systems [7, 41, 31]. [17] implements perceptual empathy in the form of a basal affect program as an autonomous reaction scheme as a route to the implementation of morality in robots. Young children and chimpanzees are equipped with this fundamental form of empathy which forms the basis for pre-social behavior [56, 26]. Imaginative empathy is much more complex and develops on the basis of perceptual empathy only. It is exhibited only in human socialisation, not in non-human primates. This form of empathy is cognitively more ambitious and is involved in more complex moral reasoning and acting [24]. We are not aware of any attempt to implement imaginative empathy in artificial systems.

Since, in a bottom-up approach, the AI system becomes a moral agent (if only in the narrowest sense of the word) one might ask whether it is likely to be more adaptable to making ethical choices in situations that are not pre-determined (which is a strong limitation of the top-down approach). Since the system learns its own ethical rules, it circumvents, to an extent (one could argue), the need to choose one particular ethical theory to implement. But this seems at least questionable. Every self-learning system must still be configured to pay attention to particular *features* of the data set, and to ignore others. Looking at the *consequences* of an action, instead of the agent’s *motivation* (for example) is such a choice of features that essentially determines the choice of moral theory. It seems difficult to judge at this point whether we can hope to create ethical-theory-agnostic AI systems, since every choice of relevant data features is already, to some extent, a choice of moral theory.

A major challenge with bottom-up approaches is it is hard to certify whether the system fulfills any requirements one might want to impose. Indeed this is a challenge for all machine learning systems. In [2] an *ethical Turing test* is proposed to tackle this issue. Under an ethical Turing test, both the AI system and an ethicist resolve the same dilemmas. The system passes the test if its choices are sufficiently similar to the ones of the ethicist. Nonetheless, this uncertainty is likely to mean bottom-up approaches are unsuitable for implementation in critical systems. This fundamental problem occurs irrespective of whether the ethical system itself or only low-level sub-systems are built in a bottom-up fashion.

## 2.4 Modular and Hybrid Approaches

Both the top-down and the bottom-up approaches come with advantages and challenges, but they also can complement each-other. A system is an entity comprised of several entities, thus in principle an AI system can be built using components that implement a top-down approach to ethical reasoning and components that implement a bottom-up approach. We are unaware of any implemented hybrid ethical reasoning system, but we can very briefly discuss some recommendations for how such a system can be created.

One approach would be to separate decision-making by the ethical prin-

ciples it involves. For example, decisions involving the possibility of human death should be made using a pre-programmed ethical policy, while decisions involving violation of autonomy can be based on ethical preferences learned through interaction with the system’s owner. Another approach would be to separate decision-making in different contexts, with bottom-up approaches being the default ethical decision-making method, while top-down approaches are implemented for certain specific pre-determined contexts. Alternatively a system can be designed so it first learns to recognize the ethical implications of its actions and then those implications can be used to follow an implemented ethical theory when choosing some specific course of action.

Implementing ethical reasoning within a system is not sufficient, we must execute such implementation in a way that allows for verification of the quality of the resulting ethical behavior. The designers and manufacturers of AI systems necessarily have to offer reasons for their users to trust the artificial ethical system, and they also need to foresee possible malfunctions and provide means to deal with them.

### 3 Specification and verification of ethical behavior

Within our society, entities that are in a position to do us harm, be it a complex machine production tool, the surgeon operating on our unconscious body, the other drivers on the highway, or a chainsaw, are subject to licensing and certification. Certification informs consumers and experts of the properties of a product, a system, or a person in a position of responsibility. Knowing that a standard has been met allows individuals to have confidence in using machinery and to trust the decisions and actions of professionals. Tools and systems are certified to operate within designated parameters, while under (well defined) proper care. Certification confirms that the manufacturer has taken all steps necessary to avoid or minimize foreseeable risks that arise in relation to the usage of the tool. Certification for persons in position of responsibility is more complex because it involves a (possibly continuous) examination to demonstrate that the certified person has the understanding and skills necessary to perform his/her duties. Typically, this involves regulations prescribing *expected* behavior — often, humans must pass an examination concerning these regulations. Once we move to an autonomous system, with no human directly in control, what are our means to ensure that a systems actually matches the relevant criteria?

In order to be confident in a system’s behavior we need to *specify* what we can expect the system to do in a particular circumstance, *verify* that the system does actually achieve this, and *validate* that our requirements are actually what the end-users want. There exist a vast range of different techniques, for example developed over many years within the field of *Software Engineering* [49]. These techniques range from the *formal*, such as proof, through *structured*, such as testing, to *informal*, such as user validation. All these approaches can, in principle, be applied across the range of autonomous systems, including robotics [20].



### 3.1 Who is the confirmation of ethical behavior for?

What constitutes an appropriate specification and verification methodology for ethical behavior depends on who is to use the results. In the case of intelligent autonomous systems at least three interested parties can be discerned: the manufacturers including developers and engineers working on developing and maintaining the systems, the end-users, owners or customers, and lastly various government and trade regulatory bodies and insurance agents. Although these three categories are the evident interested parties, this issue of interest discernment is in an open problem in its own right, and as some preliminary investigations show<sup>5</sup> finer discernment may be required.

Although those actually constructing the AI system may have an intimate knowledge of its internal workings, it is still important that developers and engineers not only have confidence in their prototypes but have techniques for highlighting where issues still remain. The technology itself should not be a black box, but should be open to maintenance and analysis, and must be flexible enough to be improved dynamically.

For end-users, customers and owners, the primary concern is that the AI system they interact with is safe and behaves ethically with respect to the ethical norms they themselves follow. *Trust* is a key issue and, in order to have trust extended to AI systems, the user needs to be informed of its range of capabilities. The future of AI systems and their proper integration within our society is subject, paradoxically, to undue levels of both optimism and pessimism in terms of the extent to which people can trust such systems. Close attention must be paid to nurturing the appropriate level of trust.

AI systems are an exciting technological development that have long been anticipated as part of the future in various works of fiction and there is the temptation to play-up their apparent capabilities, particularly by early marketing when the producers are still seeking financiers for their products. This could lead to the customers placing an unwarranted level of trust in some technology, even when adequate disclaimers and use guidelines are outlined by the manufacturer, which can in turn lead to disastrous consequences<sup>6</sup>. Such misplacement of trust is dangerous for users in the present, and may cause society to over-react in order to limit integration of technologies which given proper time to adequately develop would have been advantageous to the same society.

The appearance of trustworthiness is similarly an issue when people interact with an AI system. For example, a robot might *appear* “experienced,” “benevolent,” or “sympathetic”. Such appearances are of particular concern for AI systems that are integrated in assisted living technologies. Concerns have been raised with respect to the impact assisted living technologies can have on the elderly [46]. Similarly, it has been shown that children who interact with robots derive expectations of them and ascribe abilities to them. We need to develop an understanding of the potential long-term effects of robots on child development [33].

Trust should play a considerable role in choosing an ethical theory to implement in AI systems. The ethical theory that is easiest to implement may not

---

<sup>5</sup><http://robohub.org/should-a-carebot-bring-an-alcoholic-a-drink-poll-says-it-depends-on-who-owns-the-robot/>

<sup>6</sup><https://www.theguardian.com/technology/2016/jun/30/tesla-autopilot-death-self-driving-car-elon-musk>

necessarily be the one that is most trusted by society. This was demonstrated in the case of utilitarianism and driver-less cars [8].

It is important to note that *trust* [5] is not equal to ethics. Trust is a social construct intimately concerned with how each individual views the behavior of a robot or system. There may well be some varieties of *objective* trustworthiness, but there will remain many varieties of *subjective* trustworthiness. Many items affect users' level of trust [44], for example, the relationship between *trust* and *harm*. If you could show that robot causes no harm, would you trust it more?

Those who must regulate AI systems and their integration within society also need *confidence* in the system. In addition, the Insurance industry needs to be clear where responsibility [48] lies and so where *liability* lies. The concept of liability is likely to be complex and may be split over several actors, such as the manufacturers, the operators and the environment. *Regulation* is crucial and first steps have been taken to go beyond safety and reliability regulations [27] into considering the ethical aspects that should be taken into account [10].

Any system operating in the real world would eventually find itself in a situation in which it will malfunction and AI systems are no exception. The question is thus how certain can one be in the verified ethical behavior of an AI system and what measures can be taken to mitigate the consequences of, and learn from, the system's inevitable failure.

In terms of safety standards, the "gold standard" is currently that of aircraft autopilot software where safety is measured as the number of accidents per miles flown. We might think that, for AI systems, at least as much confidence is needed. But, is the standard too high or even achievable? There are of course, noticeable differences between aircraft and other autonomous systems. While the operational environment of an aircraft is very controlled and limited, as the aircraft must adhere to a strictly defined flying corridor, the severity of accidents is very high. *E.g.*, any malfunction in the air is certainly fatal for all of the aircraft passengers measuring in the hundreds, whereas a miscalculation on the road does not need to be fatal since cars carry fewer passengers than airplanes. It is possible that the hours of operation per accident alone is not always the best measure to assess safety of an autonomous system, but that the severity of the damage caused and the number of individuals involved in an accident should also be taken into account [39, 29, 13, 57].

### 3.2 What do we want the system to do?

A key problem is specifying *what* our expectations of an AI system are. Although this is beginning to be codified where safety is considered, for example through *robot safety* standards<sup>7</sup>, it is less clear where the ethical/moral requirements should come from and in what form should they be represented? The BS8611 standard [10], for example, does not prescribe *what* the ethical requirements should be, but maps out the issues over which ethical decisions should be considered.

An obvious route for ethical and legal requirements is through regulatory or standards bodies. These entities have the ability to set overall standards, potentially with the help of domain experts. In addition, manufacturers may well have built in specific ethical codes that go beyond (though do not contradict)

---

<sup>7</sup>See [28] for a range of robotic safety standards.

those prescribed by regulations. Finally, the user herself may wish to input her ethical preferences, ensuring that the AI acts in a way that is personally acceptable. Since there are multiple actors that need to define and refine the ethical requirements of the system, each with varying levels of technical expertise, the issue arises of how the ethical requirements are represented for the machine and for concerned actors. No one clear methodology emerges. One possibility is to have them represented in the form of a set of legal or formal rules, as argued in [45]. Another possibility is to use a set of example scenarios developed to test specific ethical choices, as in [2]. A third, but by no means final, possibility is as a statistical envelope around a large (and possibly random) set of test cases, against which the AI system must be exhaustively assessed.

### 3.3 *How do we show that the AI systems meets the expectations?*

There is a well-established body of work tackling the *Verification and Validation* (V&V) of systems, both hardware-centred and software-rich. The aim of *Verification* is to ensure that a system meets its requirements; *Formal Verification* takes this further, not only having precise formal requirements, but carrying out a comprehensive mathematical analysis of the system to ‘prove’ whether it corresponds to these formal requirements. There are many varieties of formal verification, the most popular being *model checking* [12, 6], whereby formal requirements are checked (usually automatically) against *all* possible executions of the system. Verification, via model checking, is widely used especially for the analysis of the *safety* and *reliability* of robotic systems both in terms of physical navigation [35] and in terms of internal decision-making [15]. In terms of ethical/moral verification, it seems clear that if an AI system acts by following mathematically specified rules, we can potentially formally verify its high-level behavior. Only recently, however, has the use of formal verification for *ethical* or *moral* issues begun to be addressed [14, 16].

A practical alternative to fully formal verification is to use sophisticated *coverage-driven analysis* methods, appealing to Monte-Carlo techniques and dynamic test refinement in order to systematically “cover” a wide range of practical situations. Especially where real-world interactions and devices are involved, testing is likely to be crucial. Indeed, testing for safety and reliability of robotic systems is well-established [37]. Such model-based testing is a well-developed technology but, as we move to more complex (ethical) issues sophisticated extensions may well be required. Though such approaches are typically used *before* deployment, related techniques provide a basis for run-time verification and compliance testing [42]. Testing is not as exhaustive as formal proof, but can cover many more scenarios.

*Validation* is the process of confirming that the final system has the intended behavior once it is active in its target environment, and is often concerned with satisfying external stakeholders. For example, does our system match ethical standards or legal rules set by regulators? Does our system perform acceptably from a customer point of view, and how well do users feel that it works [30]? There are many approaches to carrying out validation, typically involving the assessment of accuracy, repeatability, trust, usability, resilience, etc. All must be extended to cope with ethical and moral concerns.

It is clear that the strength and breadth of V&V research should allow us

to extend and develop this towards ethical and moral concerns. However, a number of issues remain, as follows.

- If the core software is not purely rule-based, for example involving some sub-symbolic learning procedures, then we will need a symbolic representation of the learned content if we are to carry out formal verification of the above form. One of the limitations of both formal verification and testing is likely to be in verifying learning procedures, especially where new ethical principles and preferences of behavior are learned.
- Fully formal verification is likely to be unrealistic for complete, complex systems both because of non-symbolic components (as mentioned above) and because of practical complexity limits.

However, we can formally verify *parts* of the system under particular circumstances. There are things that can be proved about core parts of *the system* and about the system's *outputs*. Consequently, formal verification techniques can provide *some* evidence. In assessing how much confidence we need in the V&V of AI system ethics, it may be possible to leave the burden of this decision to the regulator, manufacturer or end-user as appropriate. So long as a clear indication of the extent of the V&V of a system exists a user or other interested may take the decision about the risk involved in using the system. Note that we can potentially separate regulation from verification and so allow a variety of different V&V techniques to be applied.

## 4 Transparency and accountability

The choice of the relevant criteria for an AI system to be deemed ethical will eventually need to be taken by society as a whole. Therefore *transparency* is of utmost importance and thus ensuring transparency is a major challenge. To this end it is necessary to identify *what* has to be transparent to *whom*, and *how* this can be realized.

Transparency is a key requirement for ethical machines. Important attributes flow from transparency including *trust*, because it is hard to trust a machine unless you have some understanding of what it is doing and why, and *accountability*, because without transparency it becomes very difficult to understand who is responsible when a machine does not behave as we expect it to<sup>8</sup>. An ethical machine will need to be transparent to different stakeholders in different ways – each suited to that particular stakeholder. In this section we consider the transparency needs of a range of stakeholders before considering aspects of transparency common to all. This section outlines how and why transparency is important to four different groups of stakeholders: users, regulators (including accident investigators), ethicists/lawyers and society at large. Each group has different transparency needs, some of which will have to be met by allowing an AI system's ethics, and ethical logic, to be human readable, or through public engagement. Other needs will require new human-robot interfaces.

There is a relatively small literature on transparency in AI and autonomous systems. [38] proposes a theoretical framework for providing transparency in computational intelligence (CI) in order to expose the underlying reasoning

---

<sup>8</sup>Although it is important to note that transparency is not the same as accountability.

process of an agent embodying CI models. In a recent book Taylor and Kelsey [52] make the case for the importance of transparency in AI systems to an open society. For autonomous robots [59] describes early results showing that building transparency into robot action-selection can help users build a more accurate understanding of the robot. There is also no doubt that transparency is high on policy agenda: the 2016 UK Parliamentary Select Committee on Science and Technology’s final report on Robotics and AI expresses concerns over both decision making transparency and accountability and liability<sup>9</sup>. Indeed the EU’s new General Data Protection Regulation, due to take effect as law in 2018, creates a “right to explanation” such that a user will be able to ask for an explanation of an algorithmic decision that was made about them [25].

## 4.1 Transparency to the user

Although the critical importance of the human-machine interface is well understood, what is not yet clear is the extent to which an ethical machine’s ethics should be transparent to its user. It would seem to be unwise to rely on a user to discover a machine’s ethics by trial and error, but at the same time a machine that requires its user to undergo a laborious process of familiarisation may well be unworkable.

For care robots for instance it may be appropriate for the user to configure the “ethics” settings (perhaps expressing the user’s preference for more or less privacy) or, at the very least, allowing the user to choose between a small number of “preset” ethics options. There is of course always some danger that many users will rely on the default setting. What is clear is that how these options are presented to the user is very important [33]; they should for instance help and guide the user in thinking about their ‘value hierarchy’. The robot might for instance explain to the user “what would happen” in different situations and hence guide their preferences [53].

For other robot types, driverless cars for instance, the ethics settings may be fixed (perhaps by law) and therefore not user configurable. However, the need for the user to understand how the car would behave in certain situations remains critically important – especially if the car’s design (or the law) requires her to act as a safety driver and assume manual control when the autopilot cannot cope. Even for fully autonomous cars in which the user is only ever a passenger, the person with legal responsibility for the car should be aware of the car’s ethics settings. For fully autonomous cars there should still be some user interface so that the passenger can discover, or perhaps ask for help, if the vehicle become unexpectedly immobile or starts behaving erratically.

## 4.2 Transparency to regulatory bodies

It is clear that the ethics of ethical robots needs to be transparent to those responsible for (i) certifying the safety of ethical machines, and (ii) accident investigators. Both regulators and accident investigators will be working within a governance framework which includes standards and protocols. The role of the protocols is to set out how robots are certified against those standards, and – following an accident – how the accident is investigated.

---

<sup>9</sup><http://www.publications.parliament.uk/pa/cm201617/cmselect/cmsctech/145/145.pdf>

*Regulators* will need the ability to determine that a machine’s ethics comply with the appropriate standards<sup>10</sup>, and making such a determination will require those ethics to be coded and embedded into the robot in a readable way. We might imagine something like a standard Ethics Markup Language (EML – perhaps based on XML) which codes the ethics. The EML script would be embedded in the robot in a way that is accessible to the regulator, noting that the script will need to be secured to prevent attack from hackers.

*Accident Investigators.* When serious accidents happen, as they inevitably will (see Section 3.1), they will need to be investigated. To allow such investigation, data must be recorded, suggesting the need for a robot equivalent of the flight data recorder. We therefore propose an **Ethical Black Box** (EBB) – a device that records all relevant data including, crucially, internal state data on the robot’s ethical governor. Although the data stored by the EBB would be vital for investigating all aspects of an accident, including causes unrelated to the robot’s ethics, here we are interested in accidents which might have been caused by a fault or deficiency in the robot’s ethics programming. By recording the sequence of internal states of the ethical reasoning in the moments before the accident, the EBB would allow an investigator to discover exactly why the robot made an incorrect decision. Information that would be important both in determining accountability, and to make recommendations for upgrading the robot’s ethics and prevent the same accident happening again.

Specifying the EBB is beyond the scope of this paper, but it is clear that research is needed to determine what data the EBB must record, the frequency and time window of that data, and how the privacy of that data is maintained. One thing we can be sure of however is the need for an industry standard EBB (as in the aviation industry). Different EBBs and EBB standards will of course be needed for different applications, but for driverless cars for instance, a single standard EBB should be mandated. Such an EBB would itself require an industry standard, and protocols for certification, fitting and maintenance of EBBs.

### 4.3 Transparency to ethicists / lawyers

A third group of stakeholders includes lawyers, who might be required to advocate for AI systems’ owners, or on behalf of anyone who makes a claim against an AI system’s owner, or ethicists who might, for instance, be required to act as expert witnesses, in a court of law. If we consider an accident in which a robot’s ethics are implicated (see Section 4.2 above), it is clear that both lawyers and ethicists will need to understand (i) a robot’s ethics, (ii) the process the robot uses to make an ethical decision (in other words how its ethical reasoning works), and (iii) the data captured by the ethical black box. Providing this kind of transparency to lawyers and ethicists will not only be necessary, but is also likely to be challenging, as robot manufacturers may regard such details, especially (ii), as proprietary IP.

Another category of expert stakeholder includes psychologists, who might be required to either evaluate robots for their potential to cause psychological harm to the user, or as expert witnesses, in providing an investigation with an expert evaluation of the psychological harm caused to robot user(s) in a particular case.

---

<sup>10</sup>Standards Which do not yet exist.

## 4.4 Transparency to the whole of society

AI systems – and especially ethical AI systems – are a disruptive technology, with potentially significant societal and economic impact, thus an easily overlooked but important stakeholder is society as a whole. We only need to consider driverless cars and trucks to appreciate the level of potential disruption, to jobs and transport policy for instance, as already reflected in the level of public and press interest in this technology.

It is therefore very important that the ethics of ethical AI systems should be transparent to society at large, for two reasons. First, because citizens should be able to make informed judgments about the kinds of AI system they wish to have in their lives, and even more importantly those they do not want in their lives, so that they can lobby their elected representatives and ensure that government policy properly reflects those views. And second, if society is to have confidence in the ethics of a class of ethical AI system (driverless cars, for example) then it should accept a degree of collective responsibility for those ethics.

## 4.5 Technical means to bring about transparency

It is clear that the different stakeholders outlined above have very different transparency needs. Some of those needs are met through making the ethical rules and logic readable (for instance for regulators, ethicists or lawyers), but for others transparency can only be met through technical means. Here we briefly outline several approaches to meeting those needs.

- Assisted living AI systems would benefit from a “Why did you do that?” button which, when pressed, causes the robot to explain – perhaps using speech synthesized text – why it carried out the previous action. We could call the system behind this an “explanation module”. For an AI system with a fixed set of responses the explanation module should be relatively easy to implement, but for an AI system which learns its ethics such an implementation could be challenging; in either case the explanation module and its user interface would need very careful design in order to meet the needs of a non-technical user.
- An ethical AI system which makes use of simulation based internal models as part of some ethical governor (for example [58]) might allow us to go further than the “Why did you do that?” button, by making the robot’s internal simulation accessible to the user. This would enable the user to ask the robot “What would you do?” in a given situation. Clearly such a facility would need a much more sophisticated user interface than a button press, but through visualisation tools we can imagine the user watching the robot’s internal simulation running through various scenarios on a connected laptop or tablet device. Note that a similar visualisation interface would be of great value to accident investigators (Section 4.2), and expert witnesses or lawyers (section 4.3) to *play back* a robot’s internal simulation in the moments leading up to an accident, and what the alternatives open to the robot at the time might have been.
- The technical requirements for an ethical back box (EBB) were already outlined in Section 4.2 above.

## 5 Dangerous and Deliberately Unethical AI

Finally, it is important to be aware of the ways people may abuse or manipulate AI systems. As with all technology AI systems can also be deliberately abused for malice or to further one’s illegal goals [62]. While our primary concern is to contribute towards designing AI systems that behave ethically within a human society [50, 60] and promote human and animal welfare, some concern also needs to be raised about how that AI system can protect itself against abuse [63]. By abuse we, of course, do not mean mistreating the AI system in the sense in which a person or an animal can be mistreated, but taking advantage of the capabilities and opportunities offered by the AI system to commit criminal acts.

The abuse of an AI system can be achieved by hacking an existing system or by deliberately creating an unethical AI system [40, 54]. Hacking itself can be accomplished in several ways. The code of the AI system might be directly hacked. But a system can also be manipulated by interaction and such manipulation does not necessarily require technical knowledge. This is illustrated by the short-lived Tay experiment. Tay was an artificial intelligence chatter-bot released by Microsoft Corporation on March 23, 2016 and taken offline 16 hours after launch<sup>11</sup>. Tay was programmed to learn from conversation, however it took the netizens a very short time to “train” it into making morally questionable statements.

Manipulation by interaction can be accomplished both deliberately and by accident. A learning based system can be led [64] into eliciting bad conclusions through crafted case descriptions, etc. By this means one can slowly train systems away from moral behavior. As an example of accidental manipulation consider the example of children learning that driverless cars slow down in their presence, they might choose to make a game out of it. Children playing with car’s reactions might annoy passengers by causing delay; and might ultimately lead to the disabling of safeguards.

Purposeful creation of Malevolent AI can be attempted by a number of diverse agents with varying degrees of competence and success. Each such agent would bring its own goals/resources into the equation, but what is important to understand here is just how prevalent such attempts will be and how numerous such agents can be. For example, we should be concerned about: the *military* developing cyber-weapons and robot soldiers to achieve dominance; *governments* attempting to use AI to establish hegemony, control people, or take down other governments; *corporations* trying to achieve monopoly, destroying the competition through illegal means; *villains* trying to take over the world and using AI as a dominance tool; *black hats* attempting to steal information, resources or destroy cyber infrastructure targets; *doomsday cults* attempting to bring the end of the world by any means; the *depressed* looking to commit suicide by AI; *psychopaths* trying to add their name to history books in any way possible; *criminals* attempting to develop proxy systems to avoid risk and responsibility; *AI risk deniers* attempting to demonstrate that AI is not a risk factor and so ignoring caution; and even *AI safety researchers*, if unethical, attempting to justify funding and secure jobs by purposefully developing problematic AI.

The ethical and unethical behaviors of an AI system are not necessarily symmetrical. Existing systems define only a small part of the problem space

---

<sup>11</sup><http://www.bbc.com/news/technology-35890188>



[61]. Apart from ethical and unethical behavior, an AI system can also exhibit a behavior that has neither been programmed nor predicted as a particular combination of otherwise ethical rules and choices.

Lastly we must mention the potential for “cultural imperialism” when designing the ethical behavior of an AI system. With globalisation, a product’s production and consumers are diverse. What constitutes ethical behavior in one region may even be considered unethical in another. All the involved actors, the manufacturers, users and society, both on the supplier and on the demand end of the AI system need to be aware of the reality that the supplier society ethics influences the ethical behavior of the AI system, which in turn influences the ethics of the society in which the AI system operates.

## 6 Summary

Moral philosophy has a very rich history of studying how to discern right from wrong in a systematic, consistent and coherent way. Today we have a real need for a *functional* system of ethical reasoning as AI systems that function as part of our society are ready to be deployed. Building an AI system that behaves ethically is a multifaceted challenge. The questions of which ethical theory should be used to govern the AI system’s behavior has received the most of the attention. Here, we focus on the problem that comes next, after what is right or wrong for a machine to do is decided – how to implement the ethical behavior.

The problem of engineering ethical behavior is made complex because of the prime motivators for such behavior. While for humans, the motivation for behaving ethically is internal, for AI systems this motivation is external and it comes from several stakeholders. We cannot claim that the full list of these stakeholders can even be known before the AI systems are fully deployed, but we can discern between the three most evident groups of stakeholders: the manufacturers, the users and the various regulatory organs of society. Each of these stakeholders needs to play their role in deciding what is the best ethical behavior for a given AI system, but they also need to be convinced in an adequate way that the implemented behavior actually yields the desired results. A moral AI system needs to be adequately transparent and accountable to each group of stakeholders.

Unlike people, who more or less share the same “hardware” and reasoning capabilities, machines and AI systems can be built using many different approaches. The implementation of ethical reasoning will depend not only on what the stakeholders need and desire, but also on what is possible given the chosen problem-solving implementation. We discussed the two basic implementation approaches, the top-down and bottom-up approach, and identify the challenges and advantages of each.

An AI system capable of ethical behavior is necessarily a complex system. With complex systems two things are evident: that they will malfunction and that they can be used to attain criminal goals. We discuss methods of verifying that an AI system behaves as designed within specified parameters, but we also discuss how the engineering of the ethical behavior impacts available options once a system malfunctions. Lastly we discuss in broad strokes what the stakeholders need to be aware of in terms of abuse of an AI system with ethical behavior capabilities, both when that abuse is intentional and accidental.

## References

- [1] C. Allen, W. Wallach, and I. Smit. Why machine ethics? *IEEE Intelligent Systems*, 21(4):12–17, 2006.
- [2] M. Anderson and S. Leigh Anderson. Geneth: A general ethical dilemma analyzer. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, July 27 -31, 2014, Québec City, Québec, Canada.*, pages 253–261, 2014.
- [3] M. Anderson and S. Leigh Anderson. Machine ethics: Creating an ethical intelligent agent. *AI Magazine*, 28(4):15–26, 2007.
- [4] M. Anderson, S. Leigh Anderson, and Vincent Berenz. Ensuring ethical behavior from autonomous systems. In *Artificial Intelligence Applied to Assistive Technologies and Smart Environments, Papers from the 2016 AAAI Workshop, Phoenix, Arizona, USA, February 12, 2016*, 2016.
- [5] R.C. Arkin, P. Ulam, and A. R. Wagner. Moral Decision Making in Autonomous Systems: Enforcement, Moral Emotions, Dignity, Trust, and Deception. *Proceedings of the IEEE*, 100(3):571–589, 2012.
- [6] Philip J. Armstrong, Michael Goldsmith, Gavin Lowe, Joël Ouaknine, Hristina Palikareva, A. W. Roscoe, and James Worrell. Recent Developments in FDR. In *Proc. CAV*, volume 7358 of *LNCS*, pages 699–704, 2012.
- [7] M. Balconi and A. Bortolotti. Detection of the facial expression of emotion and self-report measures in empathic situations are influenced by sensorimotor circuit inhibition by low-frequency rtms. *Brain Stimulation*, 5(3):330 – 336, 2012.
- [8] J.F. s Bonnefon, A. Shariff, and I. Rahwan. The social dilemma of autonomous vehicles. *Science*, 352(6293):1573–1576, 2016.
- [9] C. Breazeal and B. Scassellati. Robots that imitate humans. *Trends in Cognitive Sciences*, 6(11):481–487, 2002.
- [10] British Standards Institution (BSI). BS 8611 Robots and Robotic Devices — Guide to the ethical design and application. <http://www.bsigroup.com>, 2016.
- [11] A. Cangelosi and M. Schlesinger. *Developmental Robotics: From Babies to Robots*. The MIT Press, 2014.
- [12] E. M. Clarke, O. Grumberg, and D. Peled. *Model Checking*. MIT Press, 1999.
- [13] Ewen Denney and Ganesh J. Pai. Automating the assembly of aviation safety cases. *IEEE Trans. Reliability*, 63(4):830–849, 2014.
- [14] L. Dennis, M. Fisher, M. Slavkovik, and M. Webster. Formal verification of ethical choices in autonomous systems. *Robotics and Autonomous Systems*, 77:1–14, 2016.

- [15] L. A. Dennis, M. Fisher, N. K. Lincoln, A. Lisitsa, and S. M. Veres. Practical Verification of Decision-Making in Agent-Based Autonomous Systems. *Automated Software Engineering*, 23(3):305–359, 2016.
- [16] L. A. Dennis, M. Fisher, and A. F. T. Winfield. Towards Verifiably Ethical Robot Behaviour. In *Proc. AAAI Workshop on AI and Ethics*, 2015.
- [17] P. Ekman. Are there basic emotions? *Psychological Review*, 99(3):550–553, 1992.
- [18] J. W. Ellington. *Translation of: Grounding for the Metaphysics of Morals: with On a Supposed Right to Lie because of Philanthropic Concerns by Kant, I. [1785]*. Hackett Publishing Company, 1993.
- [19] M. Fisher, C. List, M. Slavkovik, and A. F. T. Winfield. Engineering moral agents - from human morality to artificial morality (dagstuhl seminar 16222). *Dagstuhl Reports*, 6(5):114–137, 2016.
- [20] Michael Fisher, Louise A. Dennis, and Matthew Webster. Verifying Autonomous Systems. *ACM Communications*, 56(9):84–93, 2013.
- [21] T. Fong, I. Nourbakhsh, and K. Dautenhahn. A survey of socially interactive robots. *Robotics and Autonomous Systems*, 42(3–4):143 – 166, 2003. Socially Interactive Robots.
- [22] P. Foot. The problem of abortion and the doctrine of double effect. *Oxford Review*, 5:5–15, 1967.
- [23] T. Froese and E. A. Di Paolo. Modelling social interaction as perceptual crossing: An investigation into the dynamics of the interaction process. *Connection Science*, 22(1):43–68, March 2010.
- [24] S. Gallagher. Empathy, simulation, and narrative. *Science in Context*, 25(3):355–381, 009 2012.
- [25] B. Goodman and S. Flaxman. Eu regulations on algorithmic decision-making and a “right to explanation”. *arXiv preprint arXiv:1606.08813*, 2016.
- [26] M. Hoffman. *Empathy and Moral Development: Implications for Caring and Justice*. Cambridge University Press, 2001.
- [27] International Organization for Standardization (ISO). ISO 13482: Robots and robotic devices — Safety requirements for Personal Care Robots. <http://www.iso.org>, 2014.
- [28] International Organization for Standardization (ISO). TC299 — Robotics, 2016.
- [29] Tim P. Kelly and John A. McDermid. A systematic approach to safety case maintenance. *Rel. Eng. & Sys. Safety*, 71(3):271–284, 2001.
- [30] Hagen Lehmann, Dag Sverre Syrdal, Kerstin Dautenhahn, GertJan Gelderblom, Sandra Bedaf, and Farshid Amirabdollahian. What Should a Robot do for you? Evaluating the Needs of the Elderly in the UK. In *Proc. 6th Int. Conf. on Advances in Computer-Human Interactions*, pages 83–88, 2013.

- [31] M. J. Mataric. Getting humanoids to move and imitate. *IEEE Intelligent Systems*, 15(4):18–24, July 2000.
- [32] Andreas Matthias. Algorithmic moral control of war robots: Philosophical questions. *Law, Innovation and Technology*, 3(2):279–301, 2011.
- [33] Andreas Matthias. Robot lies in health care: when is deception morally permissible? *Kennedy Institute of Ethics Journal*, 25(2):169–162, 2015.
- [34] C. Misselhorn. Empathy with inanimate objects and the uncanny valley. *Minds and Machines*, 19(3):345, 2009.
- [35] S. Mitsch, K. Ghorbal, and A. Platzer. On Provably Safe Obstacle Avoidance for Autonomous Robotic Ground Vehicles. In *Robotics: Science and Systems IX*, 2013.
- [36] J. H. Moor. The nature, importance, and difficulty of machine ethics. *IEEE Intelligent Systems*, 21(4):18–21, July 2006.
- [37] M. Mossige, A. Gotlieb, and H. Meling. Testing Robot Controllers using Constraint Programming and Continuous Integration. *Information & Software Technology*, 57:169–185, 2015.
- [38] P. Owotoki and Mayer-Lindenberg F. Transparency of computational intelligence models. In *Research and Development in Intelligent Systems XXIII, The 26th SGAI International Conference on Innovative Techniques and Applications of Artificial Intelligence, Proceedings*, pages 387–392. Springer, 2007.
- [39] C. Patchett, M. Jump, and M. Fisher. Safety and Certification of Unmanned Air Systems. In *Engineering and Technology Reference*. 2015.
- [40] F. Pistono and R. V Yampolskiy. Unethical research: How to create a malevolent artificial intelligence. In *25th International Joint Conference on Artificial Intelligence (IJCAI-16). Ethics for Artificial Intelligence Workshop (AI-Ethics-2016)*, 2016.
- [41] G. Rizzolatti and M. Fabbri-Destro. The mirror system and its role in social cognition. *Current Opinion in Neurobiology*, 18(2):179 – 184, 2008. Cognitive neuroscience.
- [42] Grigore Rosu and Klaus Havelund. Rewriting-Based Techniques for Runtime Verification. *Automated Software Engineering*, 12(2):151–197, 2005.
- [43] S.J. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Pearson Education, 3 edition, 2015.
- [44] M. Salem, G. Lakatos, F. Amirabdollahian, and K. Dautenhahn. Would You Trust a (Faulty) Robot?: Effects of Error, Task Type and Personality on Human-Robot Cooperation and Trust. In *Proc. 10th Annual ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 141–148. ACM, 2015.
- [45] A. Saptawijaya and L. Moniz Pereira. Logic programming for modeling morality. *Logic Journal of the IGPL*, 24(4):510–525, 2016.

- [46] A. Sharkey and N. Sharkey. Granny and the robots: ethical issues in robot care for the elderly. *Ethics and Information Technology*, 14(1):27–40, 2012.
- [47] M. A. Slote. *The Ethics of Care and Empathy*. Routledge, 2007.
- [48] J. Sombetzki. Responsibility in Crisis — From the Traditional Concept of Responsibility to Systems Responsibility, 2015.
- [49] I. Sommerville. *Software Engineering*. Pearson Studium, 2001.
- [50] K. Sotala and R. V Yampolskiy. Responses to catastrophic agi risk: A survey. *Physica Scripta*, 90(1), 2015.
- [51] K. Stüeber. *Rediscovering Empathy: Agency, Folk Psychology, and the Human Sciences*. MIT Press, 2006.
- [52] R. Taylor and T. Kelsey. *Transparency and the open society: Practical lessons for effective policy*. Policy Press, Bristol UK, 2016.
- [53] A. Theodorou, R. Wortham, and Bryson J.J. Why is my robot behaving like that? designing transparency for real time inspection of autonomous robots. In *AISB Workshop on Principles of Robotics, April 2016, Sheffield UK, Proceedings*, 2016.
- [54] D. Vanderelst and A.F. Winfield. The dark side of ethical robots. *arXiv preprint arXiv:1606.02583*, 2016.
- [55] W. Wallach, C.n Allen, and I. Smit. Machine morality: Bottom-up and top-down approaches for modelling human moral faculties. *AI Soc.*, 22(4):565–582, March 2008.
- [56] F. Warneken and M. Tomasello. Varieties of altruism in children and chimpanzees. *Trends in Cognitive Sciences*, 13(9):397 – 402, 2009.
- [57] M. Webster, N. Cameron, M. Fisher, and M. Jump. Generating Certification Evidence for Autonomous Unmanned Aircraft Using Model Checking and Simulation. *Journal of Aerospace Information Systems*, 11(5):258–279, 2014.
- [58] A. F. T. Winfield, Christian Blum, and Wenguo Liu. *Towards an Ethical Robot: Internal Models, Consequences and Ethical Action Selection*, pages 85–96. Springer International Publishing, 2014.
- [59] R.H. Wortham, A. Theodorou, and Bryson J.J. What does the robot think? transparency as a fundamental design requirement for intelligent systems. In *IJCAI-2016 Ethics for Artificial Intelligence Workshop, July 2016, New York USA, Proceedings*, 2016.
- [60] R. V. Yampolskiy. Artificial superintelligence: A futuristic approach, 2015.
- [61] R. V Yampolskiy. The space of possible mind designs. In *Artificial General Intelligence: 8th International Conference, AGI 2015, AGI 2015, Berlin, Germany, July 22-25, 2015, Proceedings*, volume 9205, page 218. Springer, 2015.

- [62] R. V Yampolskiy. Taxonomy of pathways to dangerous artificial intelligence. In *Workshops at the Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- [63] R. V Yampolskiy and MS Spellchecker. Artificial intelligence safety and cybersecurity: a timeline of ai failures. *arXiv preprint arXiv:1610.07997*, 2016.
- [64] R.V. Yampolskiy. Utility function security in artificially intelligent agents. *Journal of Experimental & Theoretical Artificial Intelligence*, 26(3):373–389, 2014.